National Maritime Foundation

https://maritimeindia.org/

Operationalising Risk Assessment for Complex Maritime Systems — A Function-Oriented 'System of Systems' Framework

Tariq Ahmad

Introduction

A critical component of any disaster management plan (DMP) is the development of a risk assessment framework to identify, evaluate and prioritise disaster risks. In conventional planning approaches, risk assessment primarily supports three decision-making questions: which hazards threaten the entity, how severe those threats are, and which risks require priority mitigation. The entity under assessment may be a physical asset such as infrastructure, or a higher-order system such as an organisation, service or operational process.

Part I of this two-part series had argued that disaster risk in complex infrastructure systems cannot be understood solely through hazard-to-asset relationships. In highly interconnected systems, disruptions frequently propagate through dependencies and functional relationships rather than through direct physical damage alone. This creates a practical challenge for disaster planners: how can such functional and systemic risks be quantified in a manner that supports planning and decision-making? This article addresses that challenge by developing a risk assessment framework for operationalising the **"Function-Oriented Approach" (FOA)** within a **"System of Systems" (SoS)** model.

Traditional risk assessment begins with a hazard-centric question:

Which hazards will affect the entity?

This requires identifying plausible hazards and estimating their potential impacts, including cascading and compounding effects. The FOA, however, reframes the problem by asking an alternate question:

How does a hazard affect the critical functions and dependencies required for that entity to function?

This shift from assets to functions fundamentally changes the analytical perspective. Rather than focusing only on whether a system component fails altogether, the FOA examines whether critical functions degrade, slow down, or become unavailable. This enables a more realistic evaluation of disaster scenarios, particularly in cases where disruption emerges from degraded performance rather than complete failure.

The 2021 Texas power crisis illustrates this distinction.[1] As discussed in Part I, the crisis emerged from a combination of reduced natural gas production during extreme winter

[1] Joshua W Busby, Kaylee Shoup Baker, Morgan D Bazilian, Ashlynn Q Gilbert, Emily Grubert, Varun Rai, Joshua D Rhodes, Sarang Shidore, Caitlin A Smith, and Michael E Webber, "Cascading Risks: Understanding the 2021

conditions and a simultaneous spike in electricity demand. It is important to note that the gas extraction system did not undergo any catastrophic physical failure as a result of the blizzard (the hazard). Instead, its *operational performance* declined. That decline in performance propagated through dependent systems and contributed to widespread power outages. If the Texas grid were to have been analysed as a network of interconnected systems linked through dependencies and critical functions, the decline in upstream performance would have emerged as a significant risk factor *even in the absence of direct asset failure.*

While the FOA–SoS framework provides a means of capturing such interdependencies, it also exposes limitations in conventional risk-assessment models. Approaches commonly used in disaster planning, including those aligned with the "United Nations Office for Disaster Risk Reduction" (UNDRR) framework,[2] typically conceptualise risk as a function of hazard likelihood, exposure, vulnerability and capacity.

In such models, risk is assessed with respect to a specific hazard acting upon a specific entity. *"Exposure"* reflects the degree of interaction between the entity and the hazard, *"vulnerability"* reflects susceptibility to harm, and *"capacity"* represents the ability to reduce or absorb consequences. More generally, of course, risk is well understood to be the product of **probability** and **consequence**.[3]

Within complex systems, the limitation of such approaches lies is not necessarily in the *"scoring"* itself, but rather, in what the term *"consequence"* is recognised as representing. Unlike the case with conventional models wherein consequences are largely interpreted through direct hazard impacts on the assessed entity, under an FOA–SoS framework, consequences are ***not*** confined to direct impacts. They may propagate across interconnected systems through dependencies, interfaces and degraded functions. As a result, ***a component with low direct hazard exposure may nevertheless trigger large-scale disruption***.

This is particularly important for critical infrastructure, where vulnerability often emerges not within a single asset but across the dependency structure of the wider system. In the Texas example, the principal vulnerability was not merely within gas production assets but lay in the dependence of power generation on uninterrupted gas supply combined with inadequate resilience in downstream systems.[4] Conventional risk matrices struggle to represent this propagated vulnerability.

Accordingly, operationalising the FOA–SoS framework requires an alternative approach to risk-assessment — one capable of representing not only hazard exposure, but also functional degradation and systemic risk propagation. The remainder of this article develops such a framework for complex infrastructure sectors such as those relevant to the maritime domain, with particular focus on its applicability to disaster-management planning.

Winter Blackout in Texas", *Energy Research & Social Science* 77 (2021): 102106, https://doi.org/10.1016/j.erss.2021.102106

[2] United Nations Office for Disaster Risk Reduction (UNDRR), *The Sendai Framework Terminology on Disaster Risk Reduction*, "Disaster Risk", accessed 30 June 2026, https://www.undrr.org/terminology/disaster-risk

[3] Terje Aven, "On How to Define, Understand and Describe Risk", *Reliability Engineering & System Safety* 95, no. 6 (2010): 623–631, https://doi.org/10.1016/j.ress.2010.01.011

[4] Busby et al., "Cascading Risks: Understanding the 2021 Winter Blackout in Texas"

Existing Functional Approaches and the Ministerial Planning Challenge

Several functional risk assessment approaches have emerged to address the limitations of conventional hazard-centric risk matrices in complex systems. Among these, the *"Functional Resonance Analysis Method"* **(FRAM)**[5] focuses on how variability in operational performance propagates across functions, particularly in systems where human and organisational factors significantly influence outcomes. The *"Functional Dependency Network Analysis"* **(FDNA)**[6] approach models dependency-networks to assess how disruptions cascade through interconnected systems. The *"Critical Function Analysis"* **(CFA)**[7] approach simplifies the problem by identifying critical functions and analysing the chains of failure that may lead to systemic disruption.

Collectively, these approaches address three key analytical requirements for complex-system risk-assessment: modelling functional variability, capturing dependency-driven risk propagation and identifying critical failure pathways. These are precisely the shortcomings of conventional asset-centric risk assessment methods.

However, despite their analytical value, these methods present practical limitations for disaster management planning at the ministerial level. Some approaches, such as FRAM, primarily provide qualitative insights into systemic behaviour rather than directly producing quantitative risk outputs. Others, such as FDNA, rely upon computationally-intensive network-modelling and require specialised expertise, detailed dependency-mapping and continuous analytical effort. Such methods are highly effective for detailed engineering or operational studies, but are less suited to recurring public-sector planning cycles.

This limitation becomes particularly important in the context of disaster-management planning in India. Under the "Disaster Management Act", 2005, ministries are mandated[8] to prepare disaster management plans covering emergency measures, response coordination, control mechanisms and continuity arrangements. However, while the Act prescribes the structural contents of such plans, it does not explicitly define the strategic purpose of ministerial-level risk assessment beyond procedural compliance.

This creates a distinct planning challenge. Individual entities within a ministry — such as ports, shipyards, vessel traffic systems, and administrative offices, all of which are under the Ministry of Ports, Shipping and Waterways (MoPSW) — may each develop facility-level DMPs using risk assessment methods suited to their local operational context. Yet facility-level planning alone cannot provide a ministry-wide understanding of systemic risk. A ministerial DMP must assess how disruptions in one entity propagate across functional chains affecting logistics, trade flows, service continuity, and dependencies with other sectors.

---

[5] Erik Hollnagel, "An Application of the Functional Resonance Analysis Method (FRAM) to Risk Assessment of Organisational Change" (Sweden, 2012)

[6] Paul R. Garvey and Cesar Ariel Pinto, "Introduction to Functional Dependency Network Analysis", in *Proceedings of the Second International Symposium on Engineering Systems* (Cambridge, MA: Massachusetts Institute of Technology, 2009).

[7] Garvey and Pinto, "Introduction to Functional Dependency Network Analysis."

[8] Government of India, *The Disaster Management Act, 2005*, sec. 37, https://www.indiacode.nic.in/handle/123456789/18558?sam_handle=123456789%2F2505

This requires a unified risk-assessment framework capable of aggregating and harmonising risk across heterogeneous entities and multiple organisational hierarchies. Such a framework must satisfy four core requirements: **First**, it must integrate diverse entities including physical assets, operational processes, and organisational units. **Second**, it must remain scalable across hierarchical levels, from individual assets and functions to facilities, constituent systems, and the ministry as a whole. **Third**, it must enable harmonised comparison across similar entities operating under different local conditions, such as multiple ports. **Finally**, and most critically, it must capture propagated risk arising from functional dependencies and cascading failures.

These requirements suggest that ministerial disaster-planning does not necessarily require fully engineered simulation models, but it does require a practical risk assessment framework that preserves the core insights of functional and systemic analysis while remaining scalable, repeatable and implementable within periodic DMP planning cycles.

Consequence Dimensions of Functional Risk

To operationalise functional risk assessment within the FOA–SoS framework, the consequences of a hazard must be represented through measurable dimensions that remain meaningful across different types of entities. Conventional approaches, including those aligned with the UNDRR framework, typically assess consequences through ordinal or qualitative interpretations of exposure, vulnerability and adaptive capacity. While such scoring supports rapid comparison across hazards and entities, it provides limited interpretability regarding the nature and extent of functional degradation. In complex systems, where consequences propagate through dependencies rather than direct hazard impacts, this lack of granularity makes systemic analysis difficult.

Under the FOA–SoS framework, the consequence term remains singular but is derived from three explicit consequence dimensions: **(1)** performance loss, **(2)** damage, and **(3)** recovery period. These dimensions provide a more interpretable basis for quantification and propagation analysis.

A useful conceptual basis for this assessment is the **"Disaster Response Curve"**,[9] illustrated in **Figure 1**. The curve represents the functional state of an entity over time, following a disruptive event. Under normal conditions, the system operates at its baseline functional state. When exposed to a hazard, the system may experience a decline in functionality, followed by a recovery phase during which acceptable operations are gradually restored. The response curve therefore captures two core characteristics of disaster consequences: the severity of functional deterioration, and the duration of disruption.

[9] Charles E. Fritz, "Disaster", in *Contemporary Social Problems*, ed. Robert K. Merton and Robert A. Nisbet (New York: Harcourt, Brace and World, 1961), 651–694.

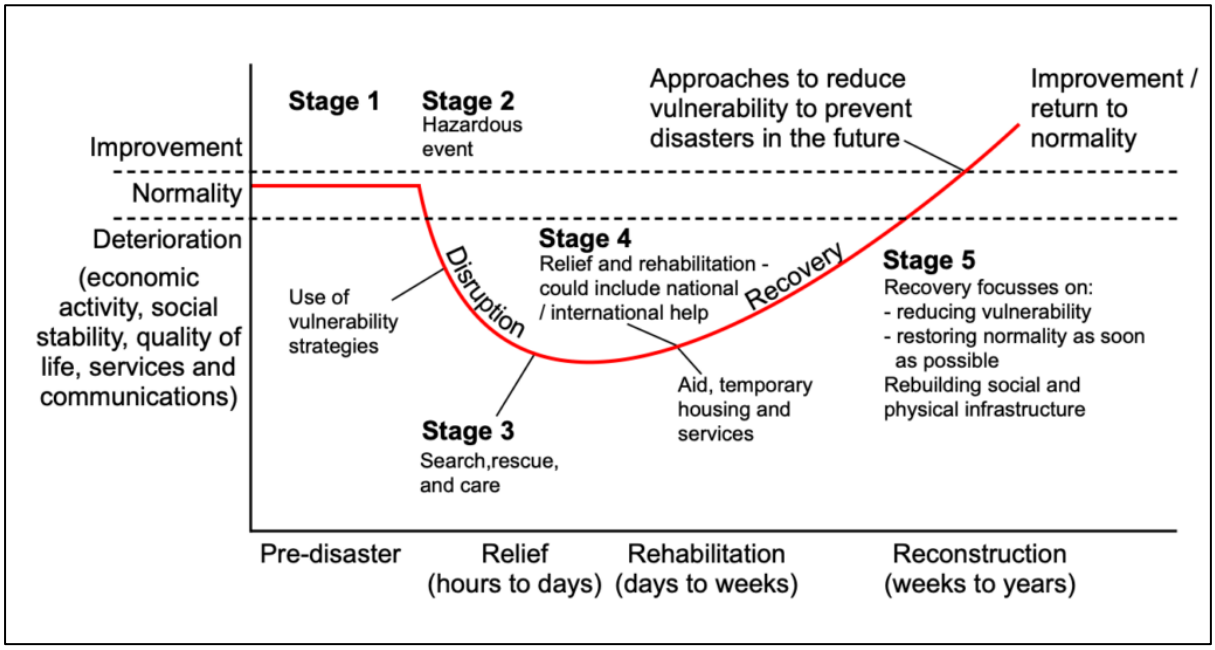

Fig 1: Disaster Response Curve

Source: Internet Geography, "Short-Term Earthquake Responses and Relief", https://www.internetgeography.net/edexcel-igcse-geography-revision/short-term-earthquake-responses-and-relief/

This perspective is particularly useful because it moves disaster assessment beyond a binary interpretation of "disaster" *versus* "no disaster". As has already been pointed out, in complex infrastructure systems, disruption frequently occurs through degraded operational states rather than complete failure. The relevant question therefore is not merely whether the system survives, but how far its performance deteriorates, what damage it sustains, and how long recovery takes.

"Performance loss" represents the reduction in functional output relative to the normal operating state. This may be *"operational"*, such as reduced cargo throughput at a port, lower vessel turnaround efficiency, degraded service delivery, etc. It may also be *"financial"*, reflecting reduced revenue generation or declining operational efficiency. Performance loss is particularly important in function-oriented assessment because it directly captures degradation in the output that downstream systems depend upon.

"Damage" represents degradation to the integrity, value or viability of the entity. This may include physical damage to infrastructure, harm to personnel, broader economic losses or reputational degradation. Some forms of damage are readily quantifiable, whereas others require structured qualitative assessment. Regardless of form, damage represents a direct consequence that affects the resilience and long-term continuity of the system.

"Recovery period" represents the time required to restore acceptable functional performance. This dimension captures not only the duration of disruption but also the resilience of the recovery process itself. In many critical systems, recovery time may be as operationally

significant as the immediate loss in performance or damage sustained, particularly where continuity requirements are stringent.

Together, these three dimensions provide a more holistic representation of hazard consequences than conventional risk matrices. They enable risk to be assessed not solely through direct physical impacts, but also through degraded performance and prolonged disruption—both of which are central to systemic failure in interconnected environments.

Within the FOA–SoS framework, these consequence dimensions are further shaped by "adaptive capacity". **"Adaptive capacity"** represents the ability of the system to absorb, withstand and recover from disruption. Importantly, it does not influence all consequence dimensions equally: some capacities reduce performance loss through redundancy, others reduce damage through protective measures, while still others shorten recovery through contingency-planning and backup systems. The challenge, therefore, lies not merely in identifying consequence dimensions, but in establishing a common basis for their quantification so that risks may be compared across heterogeneous entities and aggregated meaningfully at the system level.

Quantification Logic for Systemic Assessment

For functional risk assessment to be operational within the FOA–SoS framework, consequences must be quantified using a common and standardised basis. This is essential when risk assessments must be aggregated across heterogeneous entities such as physical infrastructure, operational processes, logistics chains, and organisational units. Any consequence metric used for systemic analysis must therefore satisfy three conditions: it must be **(1) comparable across entities**, **(2) interpretable for decision-making**, and **(3) transferable across dependency chains**.

Traditional risk-matrices often achieve comparability through ordinal scoring combined with qualitative descriptors.[10] Such approaches remain useful for rapid comparison and policy prioritisation. However, their limitations become more pronounced when the assessment must support function-oriented systemic analysis. A "Consequence Score" of 3 or 4 may indicate relative severity but provides little actionable information to dependent systems. By way of contrast, quantifying consequences through measurable outputs allows the consequence itself to become transferable. For example, a 40 per cent reduction in cargo handling capacity at a major port directly informs dependent systems—such as rail connectivity, warehousing and shipping schedules—of the extent of disruption. The metric therefore becomes not merely comparative, but usable for propagation analysis across functional dependencies.

A second limitation of conventional risk-scoring lies in combining probability and consequence into a single aggregated value too early in the assessment process. Probability and consequence serve distinct analytical purposes: probability reflects how likely a hazard scenario is to occur, whereas consequence reflects how severe disruption would be if it occurs. Combining them prematurely may suppress low-probability but high-consequence hazards, which remain highly

[10] United Nations Office for Disaster Risk Reduction (UNDRR), *Strengthening Risk Analysis for Humanitarian Planning: Integrating Disaster and Climate Risk in the Humanitarian Programme Cycle* (Geneva: United Nations Office for Disaster Risk Reduction, 2023).

relevant in critical infrastructure planning. Within the FOA–SoS framework, consequence is therefore assessed independently before being integrated with hazard probability.

To achieve systemic comparability, the three consequence dimensions—performance loss, damage and recovery period—may be expressed on a common unitless scale as percentages relative to defined reference states or thresholds. In all three dimensions, values may be capped at 100 per cent, representing the maximum tolerable consequence beyond which the entity is considered functionally failed with respect to that parameter. This cap does not imply that greater losses are impossible; rather, it reflects that consequences beyond this threshold provide limited additional planning value, since acceptable operating limits have already been exceeded.

Performance Loss

Performance loss is the most naturally suited "consequence dimension" for function-oriented quantification. Every entity operates against a baseline functional state, whether measured through operational output, service delivery or financial performance. A hazard affects the entity by reducing this output below normal levels.

As illustrated in **Figure 2**, performance loss may be represented as the percentage reduction between the normal and degraded functional states. This provides an immediately interpretable measure of functional deterioration while also enabling dependency-based analysis. If, for instance, a major port under the MoPSW experiences cyclone-induced crane disruption causing a 40 per cent reduction in cargo throughput, this degraded output can directly serve as an input for downstream risk assessments involving transport corridors, storage facilities and vessel scheduling.

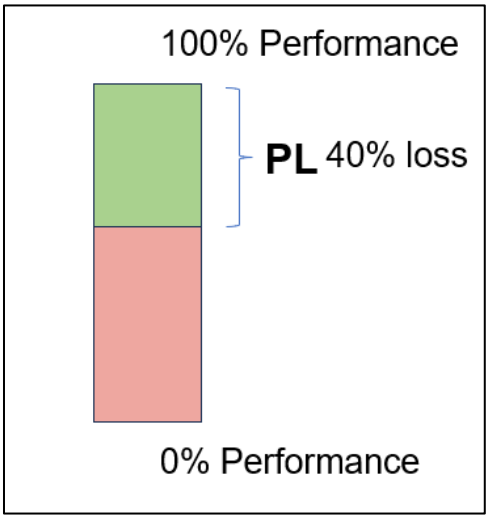

Fig 2: Performance loss (PL) as percentage of a normal reference performance
Source: Author

This principle also applies to financial performance. Reduced revenue, lower berth-utilisation, and declining throughput, may each be treated as financial performance loss, reflecting degradation in the entity's own ability to meet operational or commercial objectives. This remains conceptually distinct from broader economic damage.

Recovery Period

Unlike performance loss, "recovery period" cannot be meaningfully compared using absolute time alone. Recovery expectations vary significantly across systems: a few hours of disruption may be unacceptable for vessel traffic management or emergency communications, but tolerable for administrative functions.

Recovery period is, therefore, best assessed relative to a contextual threshold representing the maximum acceptable recovery duration, beyond which the function is considered to have effectively failed from an operational continuity perspective. As shown in **Figure 3**, actual recovery time may be expressed as a percentage of this threshold, with values capped at 100 per cent once tolerable recovery limits are exceeded.

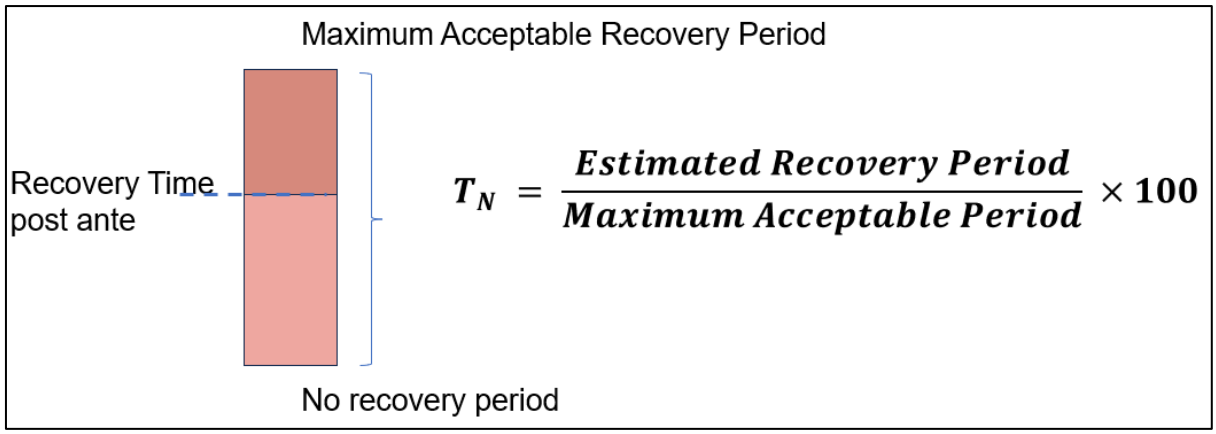

Fig 3: Recovery Period Metric as Recovery Period Ratio (T_n) with respect to a reference maximum tolerable recovery period
Source: Author

This approach preserves comparison while retaining contextual relevance. Recovery thresholds may differ depending on policy priorities, strategic importance and continuity requirements of the function being assessed.

Damage

"Damage" presents the greatest challenge for standardisation because it encompasses multiple forms of harm to the entity. Within the proposed framework, "damage" represents degradation to the integrity, value or viability of the entity. This may include physical damage to infrastructure, harm to personnel, broader economic damage and reputational damage.

Physical damage to infrastructure is generally the easiest to quantify using references such as repair cost, replacement value or degradation relative to a defined failure threshold. Similar to other consequence dimensions, damage may be represented as a percentage relative to a maximum tolerable damage threshold, capped at 100 per cent. The damage metric can be visualised in the **Figure 4**.

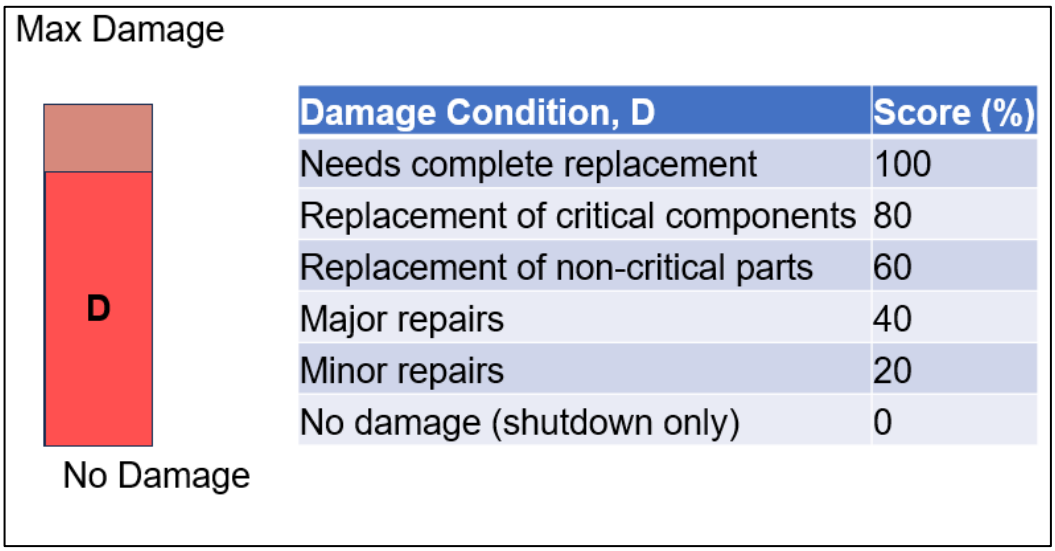

Damage Condition, D	Score (%)
Needs complete replacement	100
Replacement of critical components	80
Replacement of non-critical parts	60
Major repairs	40
Minor repairs	20
No damage (shutdown only)	0

Fig 4: Damage Metric Representation (left) with a Qualitative Metric Mapped to the Quantification Approach Proposed (right). *The values are based on damage classification scales used in structural risk assessment*
Source: Author

Damage to personnel is more complex. At the facility level, such as an individual port or terminal, serious injury or fatality may be treated as an unacceptable outcome irrespective of magnitude, making binary or near-binary assessment more appropriate. At ministerial or national scales, however, graded assessment becomes more meaningful because order-of-magnitude estimates of injuries or casualties can support comparative planning and prioritisation. Structured injury frameworks, such as the "Abbreviated Injury Scale", may provide useful reference models where graded assessment is required.

"Economic damage" should be distinguished from "financial performance loss". While financial performance loss reflects degradation in the entity's internal output or revenue, economic damage refers to broader external consequences such as supply-chain disruption, opportunity costs, contractual losses and wider market impacts arising from disruption of the entity. Reputational damage similarly affects stakeholder confidence, strategic credibility and long-term competitiveness. These forms of damage are difficult to quantify *ex ante* and are often better assessed through structured qualitative scales mapped to standardised numerical ranges.

Adaptive Capacity as a Consequence Modifier

Within the FOA–SoS framework, "adaptive capacity" remains a mitigating factor but is better understood as a *"modifier of consequences"* rather than as an independent consequence dimension. Its role is to reduce the severity of performance loss, damage and recovery period.

Importantly, "adaptive capacity" does not influence all consequence dimensions equally. Different capacities mitigate different failure modes. "**Absorptive capacities"**, such as redundancy or alternate pathways, primarily reduce performance loss. On the other hand, "**Protective capacities"**, such as structural hardening and safety systems, reduce damage, while "**Recovery capacities"** (such as contingency planning, backups and repair preparedness) reduce recovery time.

This interpretation provides significant policy-value because it enables targeted resilience planning. Rather than treating capacity as a generic variable, planners can identify specific resilience deficits within the system and prioritise interventions aimed at reducing performance loss, limiting damage or accelerating recovery. The quantified consequence dimensions, together with adaptive capacity, therefore provide the basis for aggregation into systemic risk indicators. **Figure 5** illustrates the conceptual impact of the adaptive capacity on the three parameters.

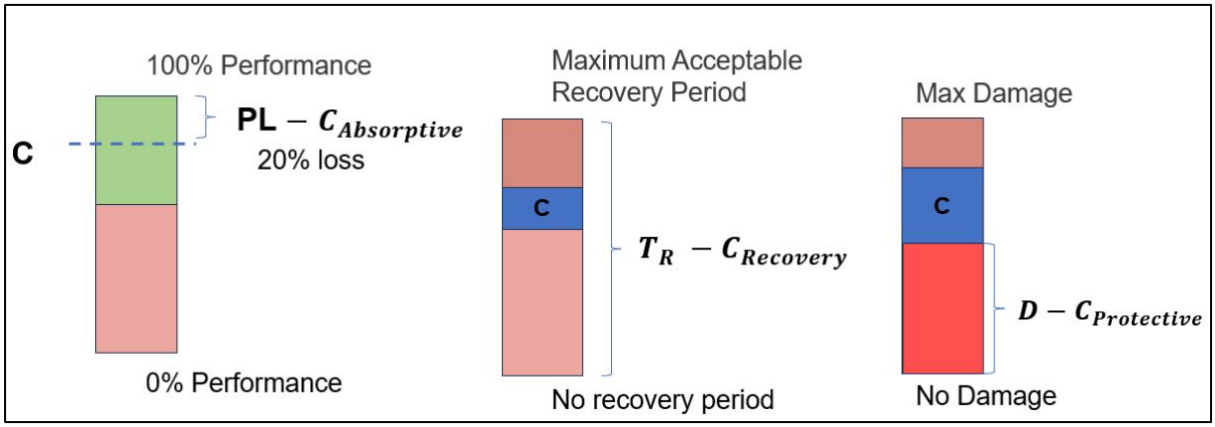

Fig 5: Capacity as a Modifier to the three Consequence Terms
Source: Author

Policy-Weighting and Functional Prioritisation

Once "consequence" dimensions have been quantified, their relative importance must be determined according to the objectives of the disaster management plan. Risk assessment is, therefore, not purely a technical exercise; it is inherently shaped by policy priorities. The severity of a consequence is meaningful only in relation to what the system is expected to preserve during and after a disaster.

Different planning contexts may assign different tolerances to the same consequence. For example, a metropolitan disaster management plan may assess and comparatively grade risk across varying magnitudes of human injuries or fatalities to support allocation of emergency-response resources. By contrast, an operational facility such as a port terminal or a shipyard may treat severe injury or loss of life as an unacceptable outcome irrespective of magnitude. Similarly, a continuously operating critical system may tolerate limited physical damage to infrastructure but have little tolerance for complete operational disruption.

Within the FOA–SoS framework, this makes the distinction between **functional degradation** and **physical damage** particularly important. Depending upon the objectives of the DMP, either consequence dimension may warrant greater priority. For the MoPSW, preserving continuity of maritime logistics and port operations may be prioritised alongside life safety and infrastructure protection. Consequently, performance loss may carry greater weight for certain critical functions, while damage-related consequences may dominate in others.

The same principle extends to sub-components within each consequence dimension. For example, damage to infrastructure, personnel, economic value, and reputation, may carry

different strategic significance depending on the planning context. Likewise, operational performance loss and financial performance loss may require different weighting depending on whether the emphasis is service continuity or commercial sustainability. These priorities may be incorporated through weighting factors at each level of assessment. Mathematically, the weights assigned to consequence dimensions and their sub-components should collectively sum to 100 per cent to preserve consistency in aggregation.

Once these policy priorities are defined, risk can be aggregated recursively across functional hierarchies to assess systemic vulnerability.

Systemic Risk Assessment

Under the FOA–SoS framework, each system is progressively disaggregated into sub-systems and further resolved into critical functions. These functions are then mapped to the assets, processes and organisational units required for their execution. Each component is assessed against relevant hazards, and the resulting consequences are aggregated to estimate risk at the functional level. **Figure 6** illustrates one such dependency-mapping.

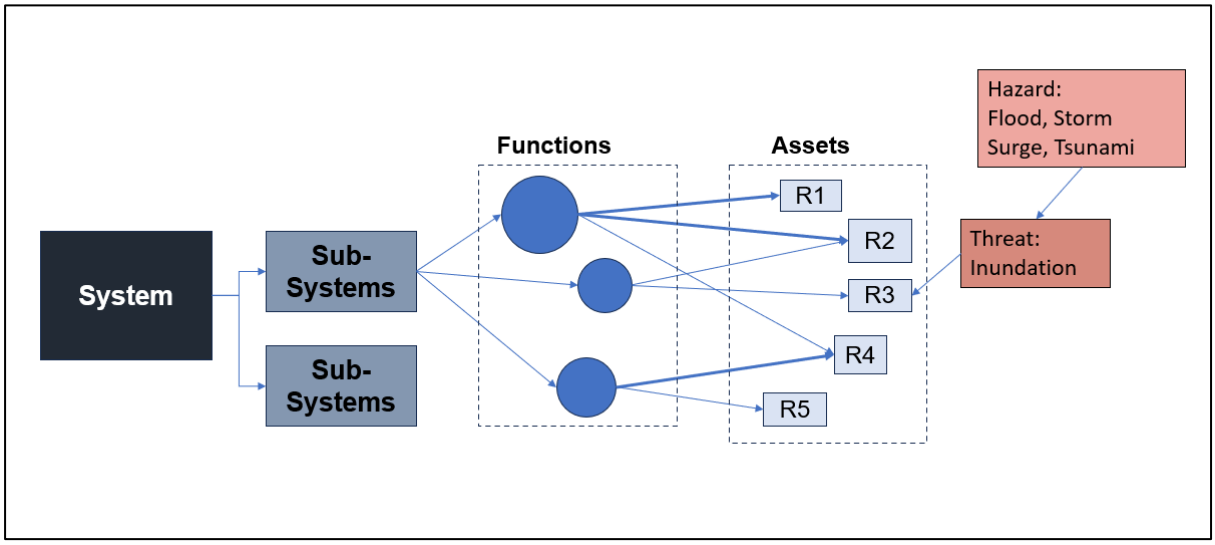

Fig 6: Risk Assessment Schema under the proposed FOA-SoS Approach
(The assets are further assigned threats impacting them and hazards leading to those threats)
Source: Author

The risk associated with a function may, therefore, be represented as a weighted aggregation of the risks of the assets and sub-functions involved in delivering that function. This naturally leads to repeated counting of risks associated with shared assets. Such repetition is not a flaw but an intentional feature of the framework. Where multiple critical functions depend on the same asset or node, the concentration of dependency increases systemic vulnerability. A system with highly centralised dependencies is inherently more fragile than one with redundancy or distributed pathways.

Once functional dependencies are mapped, systemic risk may be assessed through two complementary perspectives: "**Cumulative Risk"** and "**Critical Line of Failure Risk" (CLOFR)**.

"Cumulative Risk" represents the overall burden of risk distributed across the system. It may be understood as the weighted aggregation of risks associated with all identified functions, where weights reflect functional criticality and strategic importance. "Cumulative Risk" provides a macro-level assessment of systemic vulnerability and is particularly useful for comparing planning scenarios, assessing overall resilience and guiding broad allocation of mitigation resources. A high "cumulative risk" indicates widespread distributed vulnerability, even when no single failure pathway appears catastrophic. **Figure 7** illustrates the conceptual schema of cumulative risk.

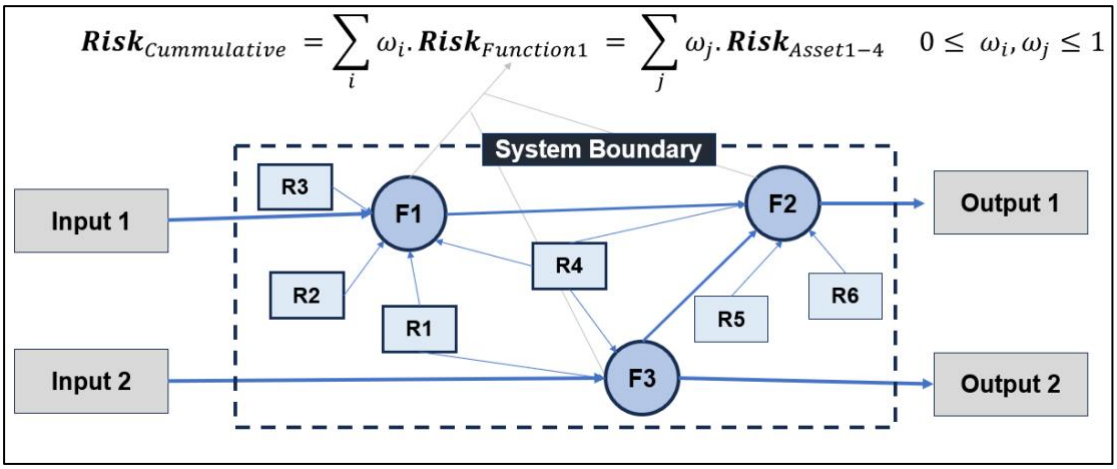

Fig 7: Conceptual schema of cumulative risk assessment. The risk is estimated for the assets (blue rectangles) associated with each function.
Source: Author

"Critical Line of Failure Risk" (CLOFR) represents the risk associated with the most critical dependency pathway whose disruption would result in failure of a critical system function. Unlike cumulative risk, CLOFR focuses not on the total burden of vulnerability but on the most consequential failure chain capable of triggering systemic breakdown. It, therefore, serves as an indicator of systemic fragility arising from critical dependencies.

Importantly, CLOFR should not be interpreted simply as the highest-risk individual asset. A high-risk asset may not necessarily lie on a critical dependency pathway. Conversely, an asset with moderate individual risk may become strategically critical if multiple essential functions depend upon it and its failure creates cascading disruption. CLOFR consequently captures vulnerability emerging from interdependence rather than isolated asset-level weakness. **Figure 8** illustrates the conceptual schema of CLOFR.

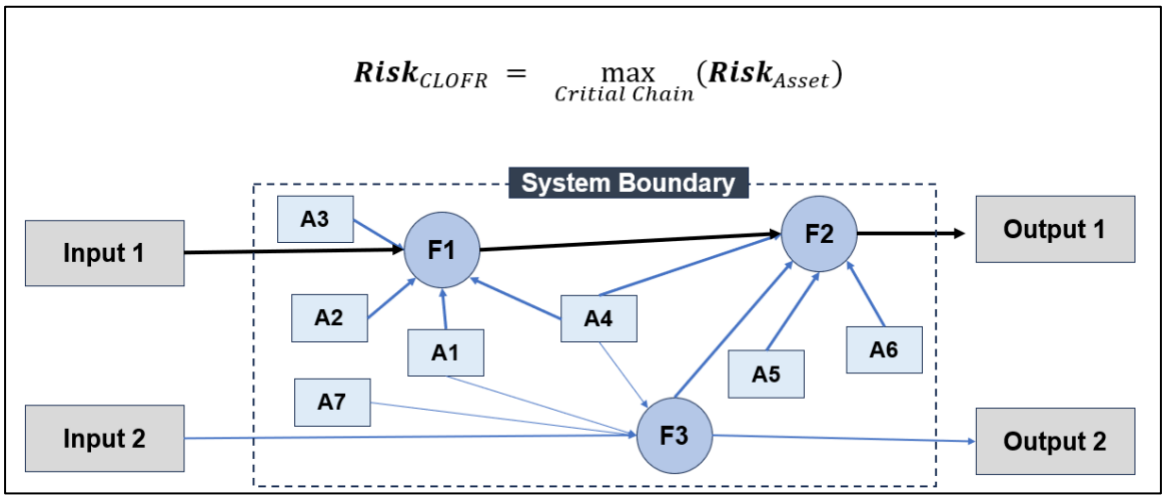

Fig 8: Conceptual Schema of CLOFR Assessment.
The risk is estimated for the asset (blue rectangles) at highest risk within an identified most critical chain (Input 1 → F1 → F2 → Output 1)
Source: Author

Taken together, Cumulative Risk and CLOFR provide two distinct but complementary decision-support outputs. Cumulative Risk helps planners understand the overall burden of vulnerability across the system, whereas CLOFR identifies critical dependency chains requiring targeted intervention. Their combined interpretation enables disaster planners to distinguish between systems requiring broad resilience enhancement and those requiring focused protection of critical functional pathways.

For ministerial-level disaster-management planning, this dual perspective is especially valuable. Large infrastructure domains such as the maritime sector rarely fail solely through isolated asset damage; systemic disruption more often emerges through cascading functional failures across interconnected systems. By combining cumulative risk and CLOFR, the FOA–SoS framework provides a practical basis for prioritising resilience investments, continuity planning and systemic risk reduction across the ministry's domain.

Conclusion and Policy Implications

This paper has examined how the "Function-Oriented Approach" (FOA) within a "System of Systems" (SoS) framework could be operationalised into a practical risk assessment methodology for disaster management planning. While Part I of the two-part paper established the conceptual limitations of asset-centric planning for complex infrastructure systems, Part II has addressed the challenge of translating that framework into measurable and comparable risk outputs for decision-making.

The central argument of this paper is that conventional risk-assessment methods remain limited when applied to highly interconnected systems, where disruption frequently emerges not through direct physical damage alone, but through degradation of functions and propagation across dependencies. In such environments, risk cannot be adequately understood solely as a hazard-

to-asset relationship. It must, instead, be assessed in terms of functional continuity, dependency-driven consequence propagation and system resilience.

The proposed framework addresses this gap by reframing consequence through three measurable dimensions—**performance loss, damage and recovery period**—which enable risk to be quantified in a manner that is both comparable across diverse entities and meaningful for systemic analysis. By incorporating adaptive capacity as a consequence-specific modifier and recursively aggregating risk across functional hierarchies, the FOA–SoS framework provides a practical basis for analysing both distributed vulnerability and critical failure pathways through **Cumulative Risk** and **Critical Line of Failure Risk (CLOFR)**.

The policy relevance of this framework becomes particularly significant in the context of ministerial disaster management planning in India. Under the National Disaster Management Authority framework established through the Disaster Management Act, 2005, ministries are mandated to prepare disaster management plans covering emergency measures, coordination arrangements and continuity planning. However, while the statutory framework specifies what a ministerial DMP should contain, it offers limited clarity regarding the strategic purpose such plans are intended to serve beyond procedural compliance. This ambiguity is important. If ministerial DMPs are treated primarily as compilations of response procedures, their utility remains largely administrative. However, if they are viewed as strategic instruments for preserving continuity of critical national functions within a ministry's domain, then systemic risk assessment becomes essential.

Accordingly, ministerial DMPs in India should move beyond facility-level hazard assessment towards function-oriented systemic risk assessment. This requires explicit identification of critical functions, mapping of interdependencies across organisational and sectoral boundaries, and incorporation of systemic resilience as a primary planning objective. For complex infrastructure sectors such as the MoPSW, where disruption propagates through logistics, transport and governance networks, such an approach is particularly necessary. As disaster risk increasingly emerges from complexity and interdependence, disaster management planning must evolve towards frameworks capable of addressing systemic failure rather than isolated asset loss.

***

About the Author

Mr Tariq Ahmad is a Research Associate at the National Maritime Foundation. His research focus is on port adaptation, disaster- and climate-change resilience, maritime spatial planning, and the blueing of the economy. His background is in architecture and spatial planning (urban & regional). He may be contacted at rsor1.nmf@gmail.com.